arXiv:1704.03966v1 [cs.CV] 13 Apr 2017

Collaborative Low-Rank Subspace Clustering

Stephen Tierney[a,*], Yi Guo[c], Junbin Gao[b]

[a]*School of Computing and Mathematics, Charles Sturt University, Bathurst, NSW 2795, Australia*
[b]*Discipline of Business Analytics, The University of Sydney Business School, The University of Sydney, NSW 2006, Australia*
[c]*Centre for Research in Mathematics, School of Computing, Engineering and Mathematics, Western Sydney University, Parramatta, NSW 2150, Australia*

Abstract

In this paper we present Collaborative Low-Rank Subspace Clustering. Given multiple observations of a phenomenon we learn a unified representation matrix. This unified matrix incorporates the features from all the observations, thus increasing the discriminative power compared with learning the representation matrix on each observation separately. Experimental evaluation shows that our method outperforms subspace clustering on separate observations and the state of the art collaborative learning algorithm.

1. Introduction

A fundamental step in machine learning and data analysis in general is to reduce the dimension of the original data to facilitate further computation by removing irrelevant information and to reduce effects such as the curse of dimensionality [1, 2]. Dimensionality reduction is most commonly performed by Principal Components Analysis (PCA) [3, 4, 5], which takes a collection of data points from their original high dimensional space and fits them to a single lower dimensional subspace. In reality however, large datasets are composed of a union or ensemble of multiple subspaces. Furthermore the subspace labels are

*Corresponding author
Email addresses: stierney@csu.edu.au (Stephen Tierney), y.guo@westernsydney.edu.au (Yi Guo), junbin.gao@sydney.edu.au (Junbin Gao)

unknown ahead of time we are presented with the two-step problem of subspace identification and fitting. The former step is known as subspace clustering.

Dimensionality reduction remains a requirement for machine learning despite increases in computational power and memory since the resolution of capture devices in all fields continues to increase. In particular dimensionality reduction is important in many Computer Vision applications due to the intrinsic high dimensional nature of of the data, since images can contain between 10 to 100 million pixels. There are many circumstances in Computer Vision under which the union of subspaces model holds and subspace clustering still remains as state of the art. The two prominent examples are identifying individual rigidly moving objects in video [6, 7, 8, 9] and identifying face images of a subject under varying illumination [10, 11]. In both these applications the subjects lie nicely on very low dimensional subspaces i.e. 4 and 9 respectively [12]. However they still present a challenge as the data can contain significant noise and corruptions. Subspace clustering has found many other computer vision applications such as image compression [13], image classification [14, 15], feature extraction [16, 17], image segmentation [18, 19] and hand-written character clustering [20, 12].

Subspace Clustering has been used outside of computer vision with success. For example it has found use in biological applications. DNA microarrays are a powerful tool in exploring gene expression and genotyping. In general the process requires counting the expression levels of genes. This process can be performed by subspace clustering [21, 22, 23, 24, 25]. At a more human scale subspace clustering has been used to automatically segment human activities from motion capture data, with the assumption that each motion class forms a subspace [26, 27, 28].

2. Preliminaries

Subspace clustering is the unsupervised classification of data points, in a set, into their respective subspace. We assume that there are p subspaces $S_1, S_2, \ldots, S_p$ with respective dimensions $\{d_i\}_{i=1}^{p}$. Both the number of sub-

spaces p and the dimension of each subspace d_i are unknown. The point set $\mathcal{X} \subset \mathbb{R}^D$ of cardinality N can be partitioned as $\mathcal{X} = \mathcal{X}_1 \cup \mathcal{X}_2 \cup \cdots \cup \mathcal{X}_p$ where each $\mathcal{X}_i$ are the data points from subspace i.

It is common to assume that each of the data points are generated using the following additive model

$$\mathbf{x} = \mathbf{a} + \mathbf{n} \tag{1}$$

where $\mathbf{x}$ is the observed data point, $\mathbf{a}$ is the true data point lying perfectly on a subspace and n is an independent and random noise term. It is rarely the case that $\mathbf{X}$ is noise or corruption free. The data is often subject to noise or corruption either at the time of capture (e.g. a digital imaging device) or during transmission (e.g. wireless communication). Two common noise assumptions for the noise type are Gaussian or Sparse noise. It is difficult to isolate the original data $\mathbf{a}$ from the noise $\mathbf{a}$ (unless the underlying structure is exactly known [29]).

Moving from an abstract set notation we can express the problem in more concrete terms by matrix notation. The observed data matrix N observed column-wise samples $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N] \in \mathbb{R}^{D \times N}$. The objective of subspace clustering is to learn the corresponding subspace label $\mathbf{l} = [l_1, l_2, \ldots, l_N] \in \mathbb{N}^N$ for each data point where $l_i \in \{1, \ldots, k\}$.

The dominant subspace clustering algorithms fall under the class of Spectral Clustering and they are our primary concern in this paper. Spectral Clustering Algorithms are defined by their used of spectral clustering to segment the data using an affinity matrix, which encodes the similarity between the data points. Spectral clustering is a popular approach to general clustering problems since it is relatively simple, computationally inexpensive, is straight forward to implement and can out perform more traditional clustering techniques [30]. They have also, in recent years, dominated the field of subspace clustering. This is because they do not increase in complexity with the number or dimension of subspaces, they are often more robust to noise or outliers and finally they provide a simple to understand work pipeline that is easily adapted and modified.

In general, spectral subspace clustering algorithms construct an affinity matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, where entry W_{ij} measures the similarity between points i and j. The hope is that $W_{ij} \gg 0$ when points i and j are in the same subspace and $W_{ij} = 0$ when they are not. Spectral clustering then constructs the segmentation of the data is obtained by applying K-means to the eigenvectors of a Laplacian matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$. In practice, Normalised Cuts (N-Cut) [31], is most often used to obtain final segmentation. N-Cut essentially follows the procedure outlined above. N-Cut has been shown to be robust in subspace segmentation tasks and is considered state of the art. In cases where N-Cut is too slow one can use approximate techniques such as the Nyström method [32].

The differences in Spectral Clustering algorithms lies in the construction of $\mathbf{W}$. One cannot directly use the distance between data points since they may be close in the ambient space but lie along different subspaces as may occur at the intersection of subspaces. Conversely data points may be distant but still in the same subspace.

Sparse Subspace Clustering (SSC) [33] is the most influential subspace clustering algorithms. This is due to its superior simplicity, accuracy, speed and robustness. SSC exploits the self-expressive property of data [12] to find the subspaces:

> *each data point in a union of subspaces can be efficiently reconstructed by a combination of other points in the data*

Which gives the relation

$$\mathbf{a}_i = \mathbf{A}\mathbf{z}_i. \tag{2}$$

where $\mathbf{z}_i$ is a vector of reconstruction coefficients for $\mathbf{x}_i$. We can then construct a model for the entire dataset as $\mathbf{A} = \mathbf{A}\mathbf{Z}$. In this unrestricted case there is no unique solution for the coefficient matrix $\mathbf{Z}$. To resolve the dilemma SSC adopts concepts from the domain of compressed sensing, namely that among the solutions to (2)

there exists a sparse solution, $\mathbf{z}_i$, whose nonzero entries correspond to data points from the same subspace as $\mathbf{a}_i$.

At the same time it is well known that a noiseless data point lying on d_i dimensional subspace S_i can be written as a linear combination of d_i other noiseless points from S_i. Therefore since it takes at minimum d_i points for the relationship to hold we can seek the sparsest representation possible i.e. a solution where the support of z is smallest.

Solving such an objective is only useful when the data is known to be noise free. As previously mentioned this is extremely unlikely in practice. To overcome this SSC assumes the data model (1) and relaxes the original constraint to yield the following objective

$$\min_{\mathbf{Z}} \lambda \|\mathbf{Z}\|_1 + \frac{1}{2}\|\mathbf{X} - \mathbf{X}\mathbf{Z}\|_F^2 \quad s.t.\ \mathrm{diag}(\mathbf{Z}) = \mathbf{0} \tag{3}$$

where λ is used to control the sparsity of $\mathbf{Z}$ and $\mathbf{X}$ is the observed data, opposed to $\mathbf{A}$, which is the latent and noise free data.

It was shown by [34] that correct subspace identification is guaranteed for SSC provided that data driven regularisation is used to setλ_i (column wise splitting of $\mathbf{Z}$). This of course is subject to further conditions such as a minimum distance between subspaces and sufficient sampling of points from each subspace.

Rather than compute the sparsest representation of each data point individually, Low-Rank Representation (LRR) [35] attempts to promote a more global affinity matrix by computing the lowest-rank representation of the set of data points. The LRR objective is as follows

$$\min_{\mathbf{Z}} \ \mathrm{rank}(\mathbf{Z}), \quad \text{s.t.} \quad \mathbf{A} = \mathbf{A}\mathbf{Z}. \tag{4}$$

The aim of the rank penalty is to create a global grouping effect that reflects the underlying subspace structure of the data. In other words, data points belonging to the same subspace should have similar coefficient patterns.

Similar to SSC, the original objective for LRR is intractable since it involves a discrete optimisation problem. Furthermore the low-rank objective has the

same problem as the ℓ_0 objective. It is only a valid objective when the data is noiseless. Instead the authors of LRR suggest a heuristic version which is the closest convex envelope: the nuclear norm. The objective then becomes

$$\min_{\mathbf{Z}} \|\mathbf{Z}\|_*, \quad \text{s.t.} \quad \mathbf{X} = \mathbf{X}\mathbf{Z} \tag{5}$$

where $\|\cdot\|_*$ is the nuclear norm of a matrix and is the sum of the singular values.

In the case where noise is present, the authors of LRR suggested a similar model to that used in SSC. However they assume that their fitting error will be only be present in a small number of columns. This results in the following objective

$$\min_{\mathbf{E},\mathbf{Z}} \lambda\|\mathbf{E}\|_{1,2} + \|\mathbf{Z}\|_*, \quad \text{s.t.} \quad \mathbf{X} = \mathbf{X}\mathbf{Z} + \mathbf{E}, \tag{6}$$

where $\|\mathbf{E}\|_{1,2} = \sum_{i=1}^{n} \|\mathbf{e_i}\|_2$ is called the $\ell_{1,2}$ norm.

3. Collaborative Subspace Clustering and Related Work

Prior work on subspace clustering has focused on the case with only a single observation of the phenomenon of interest. However in this paper paper we focus on learning a subspace segmentation from multiple observations. That is given a set of appearances $\{\mathbf{X}_1, ..., \mathbf{X}_c\}$ where c is the number of appearances we seek to jointly learn a single coefficient matrix $\mathbf{Z}$ that encodes the global coefficient structure. We illustrate this concept in Figure 1.

The motivation is two fold:

1. Combining information encoded in multiple observations has the potential to dramatically increase the discriminative power of the coefficient matrix since we can obtain a consensus, which should be more robust to outliers and noise. We demonstrate this in our later experiments.
2. Proliferation of sensors and reduced cost. Multi-sensor data is becoming increasingly common. For example in remote sensing, imagery is often captured from multiple sensors, each dedicated to particular frequency ranges

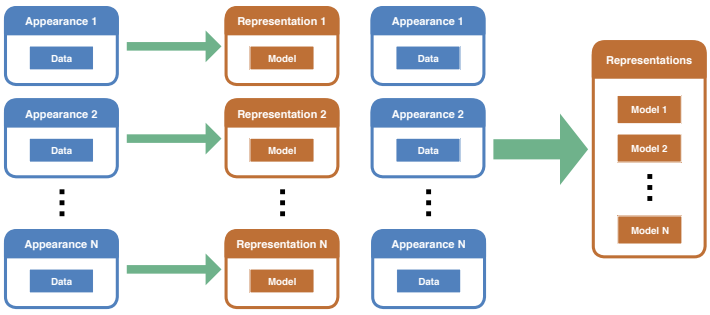

(a) Independent Subspace Clustering(b) Collaborative Subspace Clustering

Figure 1: A comparison of independent subspace clustering and collaborative subspace clustering . In independent subspace clustering (a) each appearance and affinity is considered and learnt independently. In collaborative independent subspace clustering (b) multiple affinities are learnt jointly, utilising the cross-appearance relation.

of the electromagnetic spectrum [41, 42] and in human activity recognition there has been a surge in the use of RGB-Depth cameras [43, 44]. Moreover, single observations can be turned into multiple tertiary observations. For example many image classification and recognition algorithms use many image features is used to represent the original image [19, 45, 46]. Each of these features is actually a new observation or appearance. See Figure 2 for an example.

A naive approach to this problem might be to simply perform spectral subspace clustering on each appearance and then fuse the coefficient matrices later. However the coefficient matrices could vary mildly, with little overlapping or shared structure between each $\mathbf{Z}_i$ and with vastly different magnitude of values.

The simplest solution to this problem is to simply concatenate (by row) each

(a) (b) (c)

Figure 2: An example of tertiary appearances for face image data (Extended Yale Face Database B). (a) Original Image, (b) Laplacian of Gaussian Filtered and (c) Sobel Filtered.

data matrix $\mathbf{X}_i$ to form a larger matrix e.g.

$$\mathbb{X} = \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \\ \cdots \\ \mathbf{X}_c \end{pmatrix}$$

however this has two significant drawbacks due to the heavy reliance on the Gramian matrix $\mathbf{X}^T\mathbf{X}$ in optimisation algorithms for spectral subspace clustering:

- The data in each $\mathbf{X}_i$ may not be measuring the same quantity or at the same scale. If this is the case then some rows of each $\mathbf{X}_i$ will overpower others causing a loss of information. To counter this normalisation may be used.
- The dimension of each observation ($\mathbf{X}_i$) may be different. Interestingly when the dimensions are different and normalisation is applied it is the data with the smallest dimension which tends to dominate and wash out the information encoded in the rest of the data.

With these two issues in mind it is clear that we must seek a different approach since concatenation is only suitable when each $\mathbf{X}_i$ is measuring the same quantities and is of the same dimension.

The work most closely related and which inspires our work is the so called Multi-Task Low-rank Affinity Pursuit (MLAP) [19]. MLAP adopts the low-rank

representation model for each individual observation and fuses them together by seeking sparsity consistent low-rank coefficients. The MLAP objective is the following

$$\min_{\mathbf{E},\mathbf{Z},\mathbf{A}} \sum_{i}^{c} (\|\mathbf{E}_i\|_{1,2} + \lambda_i \|\mathbf{Z}_i\|_*) + \tau \|\mathbf{A}\|_{1,2} \quad \text{s.t.} \quad \mathbf{X}_i = \mathbf{X}_i \mathbf{Z}_i + \mathbf{E}_i, \tag{7}$$

where $\mathbf{X}_i$, $\mathbf{E}_i$ and $\mathbf{Z}_i$ are observation data, the fitting error and coefficient matrices for the ith observation respectively, c is the number of observations, the $\ell_{1,2}$ norm is defined as $\|\mathbf{E}\|_{1,2} = \sum_{i=1}^{n} \|\mathbf{e_i}\|_2$, $\mathbf{A} = [\text{vec}(\mathbf{Z}_1), \ldots, \text{vec}(\mathbf{Z}_c)]^T$ and vec is the vectorisation operator such that $\text{vec}(\mathbf{C}) = \mathbf{C}^{m \times n} \mapsto \mathbf{c}^{mn}$.

However the sparsity consistency penalty is rather unsatisfactory since it will not ensure that every $\mathbf{Z}_i$ is actually similar. To demonstrate this we examine the proximal $\ell_{1,2}$ thresholding step in MLAP, which can be written as

$$\underset{\mathbf{A}}{\text{argmin}} \ \tau \|\mathbf{A}\|_{1,2} + \frac{1}{2} \|\mathbf{A} - \mathbf{V}\|_F^2 \tag{8}$$

where $\mathbf{V}$ is a placeholder for the variables involved in this step. The solution to this problem is given by

$$\mathbf{A}_i = \begin{cases} \dfrac{\|\mathbf{v}_i\|_2 - \tau}{\|\mathbf{v}_i\|_2} \mathbf{v}_i & \text{if } \|\mathbf{v}_i\|_2 > \tau \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

where $\mathbf{A}_i$ and $\mathbf{v}_i$ are columns of $\mathbf{A}$ and $\mathbf{V}$ respectively. Let

$$\mathbf{V} = \begin{pmatrix} 0.0010 & 1.0000 & 1.0000 \\ 0.0015 & 0.1000 & 1.0100 \end{pmatrix}$$

and $\tau = 0.1$ then the solution is

$$\mathbf{A} = \begin{pmatrix} 0 & 0.9005 & 0.9296 \\ 0 & 0.0900 & 0.9389 \end{pmatrix}$$

Here we have presented three possible scenarios in the three columns. From left to right: small magnitudes with a small difference, large and small magnitudes, and large magnitudes with a small difference. In the first and last columns

where the row difference is small the row difference is further narrowed. Likewise with the middle column, however the difference is not decreased by an appreciable amount. In other words the $\ell_{1,2}$ penalty appears to work best with values that are already close. But as previously mentioned this is unlikely to be the case as different appearances can produce vastly different coefficient matrices. In other words the $\ell_{1,2}$ penalty can fail to inductively transfer structure between the coefficient matrices. Furthermore the $\ell_{1,2}$ penalty suppresses the numeric values in $\mathbf{V}$. This suppression can lead to an increase in the fitting error $\mathbf{E}_i$ and reduce clustering accuracy.

After solving the MLAP objective (7) the authors then obtain a unified coefficient matrix $\mathbb{A} \in \mathbb{R}^{N \times N}$ by

$$\mathbb{W}_{ij} = \sqrt{\sum_{k=1}^{c} (\mathbf{Z}_k)_{ij}^2} + \sqrt{\sum_{k=1}^{c} (\mathbf{Z}_k)_{ji}^2}. \tag{10}$$

Noting that $\sqrt{\sum_k^c (\mathbf{Z}_k)_{ij}^2}$ is the ℓ_2 norm of all the coefficients for element at position i, j.

4. Collaborative Low-Rank Subspace Clustering

Rather than enforce that the magnitudes of the coefficient matrices $\mathbf{Z}_i$ are similar, as in MLAP, we seek only that the coefficient pattern is similar. Recall that $\mathbf{A} = [\text{vec}(\mathbf{Z}_1), \ldots, \text{vec}(\mathbf{Z}_c)]^T$. Let us consider an ideal case where the solution coefficient pattern $\mathbf{Z}_i$'s are similar i.e. they are the scaled by a constant factor from each other. In other words they are linearly dependant and thus $\text{rank}(\mathbf{A}) = 1$. Starting from this ideal case provides our motivation. Therefore we seek to solve the following objective function

$$\begin{aligned} &\min_{\mathbf{E},\mathbf{Z}} \sum_i \|\mathbf{E}_i\|_F^2 + \lambda_i \|\mathbf{Z}_i\|_*, \\ &\text{s.t.} \quad \mathbf{X}_i = \mathbf{X}_i \mathbf{Z}_i + \mathbf{E}_i, \ \text{rank}(\mathbf{A}) = 1. \end{aligned} \tag{11}$$

Unfortunately such a strict rank condition means the objective is no longer convex, which makes optimisation difficult. We can relax the constraint into

Algorithm 1 cLRSC

Require: $\{\mathbf{X}^{D\times N}\}$ - observed data, c - number of subspaces

1: Obtain coefficients $\mathbf{Z}_i$ by solving (13)

2: Form the similarity graph $\mathbb{W}_{ij} = \sqrt{\sum_{k=1}^{c}(\mathbf{Z}_k)_{ij}^2} + \sqrt{\sum_{k=1}^{c}(\mathbf{Z}_k)_{ji}^2}$.

3: Apply N-Cut to $\mathbf{W}$ to partition the data into c subspaces

4: **return** Subspace label indices for $\{S_i\}_{i=1}^{c}$

the objective function since a sufficiently large τ can simulate the $\text{rank}(\mathbf{A}) = 1$ constraint

$$\min_{\mathbf{E},\mathbf{Z},\mathbf{A}} \sum_i (\|\mathbf{E}_i\|_F^2 + \lambda_i \|\mathbf{Z}_i\|_*) + \tau \text{ rank}(\mathbf{A}) \qquad (12)$$
$$\text{s.t.} \quad \mathbf{X}_i = \mathbf{X}_i\mathbf{Z}_i + \mathbf{E}_i.$$

However this objective is NP-hard. Furthermore the rank constraint may be too strict. Enforcing that each $\mathbf{Z}_i$ is an exact linear combination of the others may force $\mathbf{E}_i$ to absorb a significant amount of error. These issues lead us to our actual final objective function for LRCSC (Low-Rank Collaborative Subspace Clustering)

$$\min_{\mathbf{E},\mathbf{Z},\mathbf{A}} \sum_i (\|\mathbf{E}_i\|_F^2 + \lambda_i \|\mathbf{Z}_i\|_*) + \tau \|\mathbf{A}\|_* \qquad (13)$$
$$\text{s.t.} \quad \mathbf{X}_i = \mathbf{X}_i\mathbf{Z}_i + \mathbf{E}_i$$

We now have a feasible objective function that can simulate the rank constraint. Just as τ can simulate the rank one constraint in (12), τ in (13) can also simulate this constraint if it is sufficiently large [47]. If a unified coefficient matrix is required we take the approach described in (10).

We outline the full Collaborative Low-Rank Subspace Clustering (cLRSC) algorithm in Algorithm (1). Compared with MLAP, cLRSC will enforce coefficients in each affinity matrix to have a similar pattern rather than similar magnitudes. We later show through experimental evaluation that cLRSC outperforms MLAP consistently.

4.1. Optimisation

To solve (13) we take the path provided by LADMPSAP (Linearized Alternating Direction Method with Parallel Spliting and Adaptive Penalty) [48]. An overview of the entire algorithm can be found in Algorithm 2. Here we expand on solving the primary variables. First introduce an auxiliary constraint to explicitly enforce the relationship between $\mathbf{Z}_i$ and $\mathbf{A}$

$$\begin{aligned} \min_{\mathbf{E},\mathbf{Z},\mathbf{A}} \sum_{i=1}^{c}(\|\mathbf{E}_i\|_F^2 + \lambda_i\|\mathbf{Z}_i\|_*) + \tau\|\mathbf{A}\|_* \\ \text{s.t.} \quad \mathbf{X}_i = \mathbf{X}_i\mathbf{Z}_i + \mathbf{E}_i \\ \text{vec}(\mathbf{Z}_i)^T = \mathbf{A}_{(i,:)} \end{aligned} \tag{14}$$

where $\mathbf{A}_{(i,:)}$ is the ith row of $\mathbf{A}$. Next we relax the constraints and form the Augmented Lagrangian

$$\begin{aligned} &\min_{\mathbf{E}_i,\mathbf{Z}_i,\mathbf{A}} \sum_{i=1}^{c}(\|\mathbf{E}_i\|_F^2 + \lambda_i\|\mathbf{Z}_i\|_*) + \tau\|\mathbf{A}\|_* \\ &+ \sum_{i=1}^{c}(\langle\mathbf{Y}_i, \mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i\rangle + \frac{\mu}{2}\|\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i\|_F^2 \\ &+ \langle\mathbf{w}_i, \text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)}\rangle + \frac{\mu}{2}\|\text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)}\|_F^2) \end{aligned} \tag{15}$$

Each $\mathbf{Z}_i$ can be updated in parallel by

$$\begin{aligned} \min_{\mathbf{Z}_i} \sum_{i=1}^{c}(\lambda_i\|\mathbf{Z}_i\|_* + \frac{\mu}{2}\|\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i + \frac{1}{\mu}\mathbf{Y}_i\|_F^2 \\ + \frac{\mu}{2}\|\text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)} + \frac{1}{\mu}\mathbf{w}_i\|_F^2) \end{aligned}$$

To find a closed form solution we linearise the second and third terms

$$\min_{\mathbf{Z}_i} \sum_{i=1}^{c}(\lambda_i\|\mathbf{Z}_i\|_* + \frac{\rho}{2}\|\mathbf{Z}_i - (\mathbf{Z}_i^t - \frac{1}{\rho}\partial G(\mathbf{Z}_i^t))\|_F^2)$$

where $G = \frac{\mu}{2}\|\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i + \frac{1}{\mu}\mathbf{Y}_i\|_F^2 + \frac{\mu}{2}\|\text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)} + \frac{1}{\mu}\mathbf{w}_i\|_F^2$ and $\partial G = -\mu\mathbf{X}^T(\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i + \frac{1}{\mu}\mathbf{Y}_i) + \mu(\text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)} + \frac{1}{\mu}\mathbf{w}_i)$. Note that the second half of ∂G needs to be reshaped into the same size as $\mathbf{Z}_i$. The solution

to this problem is given by

$$\mathbf{Z}_i = \mathbf{U} S_{\frac{\lambda_i}{\rho}}(\mathbf{\Sigma}) \mathbf{V}^T, \tag{16}$$

$$S_\beta(\mathbf{\Sigma}) = \text{diag}(\{\max(\sigma_i - \beta, 0)\}),$$

where $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is the Singular Value Decomposition (SVD) of $\mathbf{Z}_i^t - \frac{1}{\rho}\partial G(\mathbf{Z}_i^t)$ and σ_i are the singular values, see [49][50] for details.

Each $\mathbf{E}_i$ can be updated in parallel by

$$\min_{\mathbf{E}_i} \sum_{i=1}^{c} (\|\mathbf{E}_i\|_F^2 + \frac{\mu}{2}\|\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i + \frac{1}{\mu}\mathbf{Y}_i\|_F^2)$$

$$\mathbf{E}_i = \frac{\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \frac{1}{\mu}\mathbf{Y}_i}{\frac{1}{\mu^k} + 1} \tag{17}$$

$\mathbf{A}$ can be updated in parallel by

$$\min_{\mathbf{A}} \tau\|\mathbf{A}\|_* + \sum_{i=1}^{c} (\frac{\mu}{2}\|\text{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)} + \frac{1}{\mu}\mathbf{w}_i\|_F^2)$$

which we rearrange to

$$\min_{\mathbf{A}} \frac{\tau}{\mu}\|\mathbf{A}\|_* + \frac{1}{2}\|\mathbf{A} - \mathbf{B}\|_F^2$$

where $\mathbf{B}_{(i,:)} = \text{vec}(\mathbf{Z}_i)^T + \frac{1}{\mu}\mathbf{w}_i$ and the final solution is given by

$$\mathbf{A} = \mathbf{U} S_{\frac{\tau}{\mu}}(\mathbf{\Sigma}) \mathbf{V}^T. \tag{18}$$

where $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ is the SVD of $\mathbf{B}$.

5. Experimental Evaluation

In this section we evaluate cLRSC for subspace clustering. First we perform a synthetic experiment with randomly generated appearance data generated. Next we perform a semi-synthetic experiment generated from a library of real data. Finally we perform an experiment on real world data in the form of face images. Parameters were fixed for each experiment and tuned for best performance.

Algorithm 2 Solving (13) **by LADMPSAP**

Require: $\{\mathbf{X}_i, ...\mathbf{X}_c\}$, $\{\lambda_i, ...\lambda_c\}$, τ, μ, $\mu^{\max} >> \mu$, $\rho > \|\mathbf{X}\|^2$ and $\epsilon_1, \epsilon_2 > 0$.

1: Initialise $\{\mathbf{Z}_i, ...\mathbf{Z}_c\} = \mathbf{0}$, $\{\mathbf{E}_i, ...\mathbf{E}_c\} = \mathbf{0}$, $\mathbf{A} = \mathbf{0}$, $\{\mathbf{Y}_i, ...\mathbf{Y}_c\} = \mathbf{0}$, $\{\mathbf{w}_i, ...\mathbf{w}_c\} = \mathbf{0}$

2: **while** not converged **do**

3: Update each $\mathbf{Z}_i^{k+1}$ by (16)

4: Update each $\mathbf{E}_i^{k+1}$ by (17)

5: Update $\mathbf{A}^{k+1}$ by (18)

6: Set m

$$m = \frac{\mu^k\sqrt{\rho}}{\|\mathbf{X}\|_F}\max\big\{\|\mathbf{A}^{k+1} - \mathbf{A}^k\|_F, \|\mathbf{Z}_i^{k+1} - \mathbf{Z}_i^k\|_F, \|\mathbf{E}_i^{k+1} - \mathbf{E}_i^k\|, \ldots, \|\mathbf{Z}_c^{k+1} - \mathbf{Z}_c^k\|_F, \|\mathbf{E}_c^{k+1} - \mathbf{E}_c^k\|)\big\}$$

7: Check stopping criteria for each i to c

$$\frac{\|\mathbf{X}\mathbf{Z}_i^{k+1} - \mathbf{X} + \mathbf{E}_i^{k+1}\|_F}{\|\mathbf{X}_i\|_F} < \epsilon_1;$$
$$\frac{\|\mathbf{A}_{(i,:)}^{k+1} - \mathrm{vec}(\mathbf{Z}_i^{k+1})^T\|_F}{\|\mathbf{X}_i\|_F} < \epsilon_1;$$
$$m < \epsilon_2$$

8: Update each $\mathbf{Y}_i$

$\mathbf{Y}_i^{k+1} = \mathbf{Y}_i^k + \mu(\mathbf{X}_i - \mathbf{X}_i\mathbf{Z}_i - \mathbf{E}_i)$

9: Update each $\mathbf{w}_i$

$\mathbf{w}_i^{k+1} = \mathbf{w}_i^k + \mu(\mathrm{vec}(\mathbf{Z}_i)^T - \mathbf{A}_{(i,:)})$

10: Update γ

$$\gamma_1 = \begin{cases} \gamma^0 & \text{if } m < \epsilon_2 \\ 1 & \text{otherwise,} \end{cases}$$

11: $\mu^{k+1} = \min(\mu_{\max}, \gamma\mu^k)$

12: **end while**

In an effort to maximise transparency and repeatability, all MATLAB code and data used for these experiments and those in Section **??** can be found online at `https://github.com/sjtrny/collab_lrsc`.

To compare clustering accuracy we use the subspace clustering accuracy (SCA) metric [12], which is defined as

$$\text{SCA} = 100 - \frac{\text{num. of misclassified points} \times 100}{\text{total num. of points}} \tag{19}$$

Essentially it is a percentage of possible accuracy i.e. 100 is perfect clustering while 0 has no matched clusters. However due to ambiguity such a low score is not possible. In cases where we inject extra noise we report the level of noise using Peak Signal-to-Noise Ratio (PSNR) defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{s^2}{\frac{1}{mn} \sum_i^m \sum_j^n (A_{ij} - X_{ij})^2} \right) \tag{20}$$

where $\mathbf{X} = \mathbf{A} + \mathbf{N}$ is the noisy data and s is the maximum possible value of an element of $\mathbf{A}$. Decreasing values of PSNR indicate increasing amounts of noise.

5.1. Synthetic Data

In this section we compare clustering accuracy from a randomly generated dataset. Similar to [35] we construct 5 subspaces $\{S_i\}_{i=1}^5$ whose bases $\{\mathbf{U}_i\}_{i=1}^5$ are computed by $\mathbf{U}_{i+1} = \mathbf{T}\mathbf{U}_i, 1 \leq i \leq 4$, where $\mathbf{T}$ is a random rotation matrix and $\mathbf{U}_1$ is a random orthonormal basis of dimension 100×4. In other words each basis is a random rotation away from the previous basis and the dimension of each subspace is 4. 20 data points are sampled from each subspace by $\mathbf{X}_i = \mathbf{U}_i \mathbf{Q}_i$ with $\mathbf{Q}_i$ being a 4×20 random iid $\mathcal{N}(0, 1)$ matrix. The data is concatenated $\mathbb{X} = [\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_5]$. This process is repeated again yielding two appearances $\{\mathbb{X}_1, \mathbb{X}_2\}$.

We perform subspace clustering with LRR on the individual appearances, MLAP and cLRSC on both observations. Normalized cuts [31] is used for final segmentation. To obtain accurate results we repeat this experiment 50 times. Further to test robustness we repeat again for various levels of corruption by Gaussian noise. Results are reported in Figure 4. cLRSC (ours) outperforms

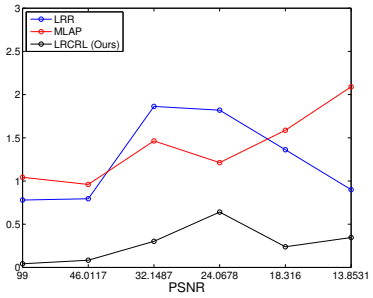

Figure 3: A comparison of mean differences between each $\mathbf{Z}_i$ for the semi-synthetic dataset.

both MLAP and LRR at all noise levels. This demonstrates our earlier hypothesis that collaborative learning of can increase the discriminative power of coefficient matrices. Further cLRSC shows robustness even in cases of extreme noise. For example Figure 4 (d) the minimum SCA at PSNR 32 is 94%. In contrast minimum SCA of MLAP for this case is 40%.

5.2. Semi-Synthetic Data

We assemble semi-synthetic data from a library of pure infrared hyper spectral mineral data. We randomly take 5 pure spectra samples from the library such that $\mathbf{A}_i = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_5] \in \mathbb{R}^{321\times 5}$. Next we combine these samples into a single synthetic sample using uniform random weights $\mathbf{w}_i \in \mathbb{R}^5$ such that $\mathbf{x}_i \in \mathbb{R}^{321} = \mathbf{A}_i\mathbf{w}_i$. We then repeat $\mathbf{x}_i$ 10 times column-wise giving $\mathbf{X}_i \in \mathbb{R}^{321\times 10}$. We repeat this process 5 times and combine all $\mathbf{X}_i$ to create our artificial data $\mathbb{X} \in \mathbb{R}^{321\times 50} = [\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_5]$. The entire process is repeated again yielding two observations of the same phenomenon $\{\mathbb{X}_1, \mathbb{X}_2\}$.

We perform subspace clustering with LRR on the individual appearances, MLAP and cLRSC on both observations. Normalized cuts [31] is used for final segmentation. To obtain accurate results we repeat this experiment 50 times. Further to test robustness we repeat again for various levels of corruption by Gaussian noise. Results are reported in Figure 5. Like the previous synthetic experiment cLRSC (ours) outperforms both MLAP and LRR at all noise levels and is far more robust at extreme noise levels.

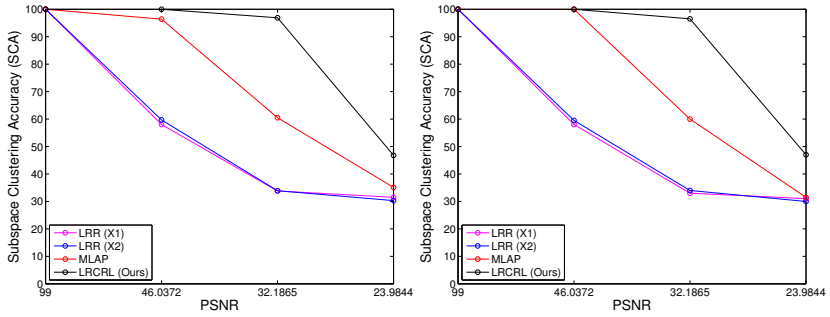

(a) Mean SCA (b) Median SCA

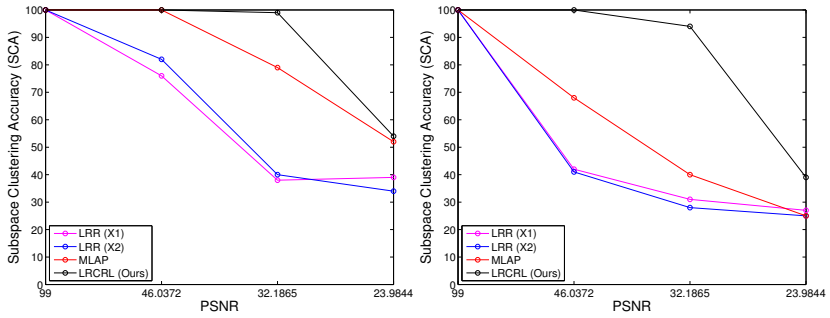

(c) Maximum SCA (d) Minimum SCA

Figure 4: Results for the synthetic data segmentation experiment with various magnitudes of Gaussian noise. cLRSC (ours) outperforms both MLAP and LRR at all noise levels.

In Figure 3 we compare how close the coefficient matrices $\mathbf{Z}_1$ and $\mathbf{Z}_2$ are for cLRSC, MLAP and LRR. We measure the mean difference through the Frobenius norm i.e. $\frac{\sum_1^{50} \|\mathbf{Z}_1^i - \mathbf{Z}_2^i\|_F}{50}$. From the results we can see that the coefficient matrices are much closer under cLRSC than MLAP and LRR. Note that in most cases the difference between MLAP coefficient matrices is actually larger than those from LRR on individual appearance data. This highlights our earlier point that MLAP cannot jointly learn coefficient matrices.

5.3. Real Data

The aim of this experiment is to segment unique human subjects from a set of face images. We draw our data from the Extended Yale Face Database B [11].

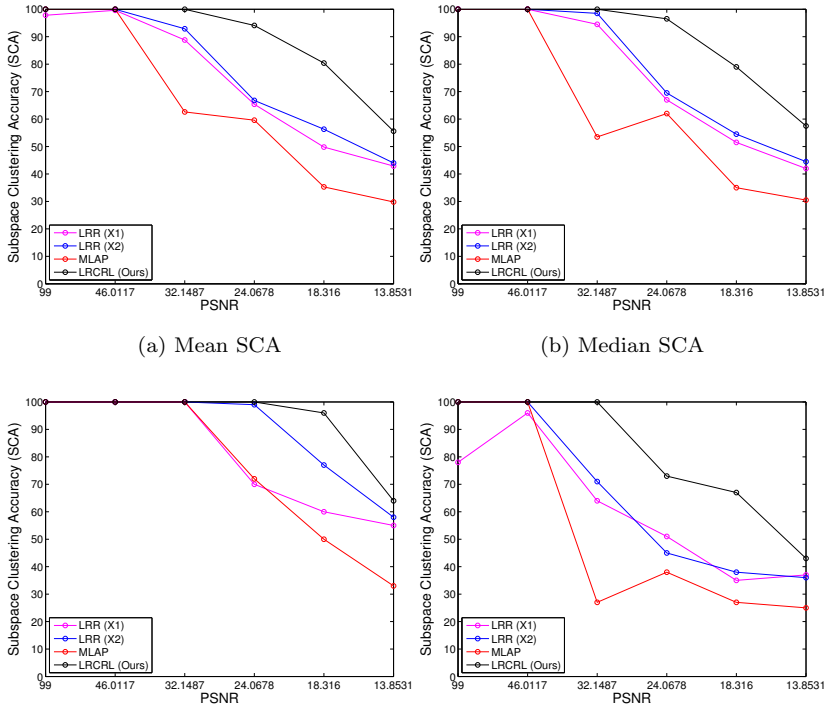

(a) Mean SCA (b) Median SCA

(c) Maximum SCA (d) Minimum SCA

Figure 5: Results for the semi-synthetic data segmentation experiment with various magnitudes of Gaussian noise. cLRSC (ours) outperforms both MLAP and LRR at all noise levels.

The dataset consists of approximately of 38 subjects under varying illumination (see Figure 2). The test consisted of randomly selecting 5 subjects and their corresponding exemplar image (minimum shadows and corruption). These images were then resampled to 42×48 pixels and then repeated 20 times column wise and finally concatenated together to yield $\mathbb{X}_0 \in \mathbb{R}^{2016 \times 100} = [\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_5]$ where $\mathbf{X}_i \in \mathbb{R}^{2016 \times 20}$.

From this data we generated two tertiary appearances which were used for subspace clustering. The first was a Laplacian of Gaussian filtered version and the second was a Sobel edge-emphasising filtered version yielding two appearances $\{\mathbb{X}_1, \mathbb{X}_2\}$. For an example of the appearance data used pleas see Figure 2. We repeated tests 50 times for various levels of corruption with Gaussian noise.

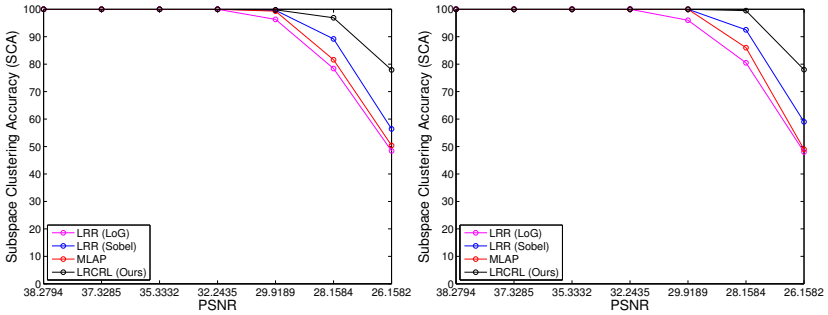

(a) Mean SCA (b) Median SCA

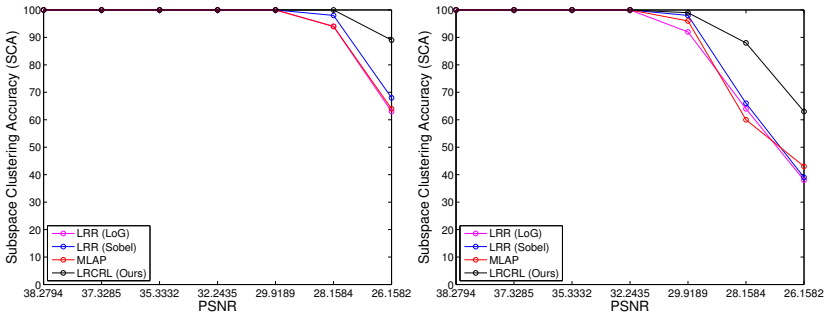

(c) Maximum SCA (d) Minimum SCA

Figure 6: Results for the real data (Extended Yale Face Database B) segmentation experiment with various magnitudes of Gaussian noise. cLRSC (ours) outperforms both MLAP and LRR at all noise levels.

Results can be found in Figure 6. cLRSC shows an improvement in clustering accuracy over MLAP and LRR on individual appearances.

6. Conclusion

We presented and evaluated a novel framework for collaborative low-rank subspace clustering (cLRSC). This framework exploits the representation encoded in multiple appearances of a single phenomenon to increase discriminative power. This framework outperforms the prior state of the art in this field (LRR and MLAP). Further it is extremely robust to noise. In the future we hope to address two areas where cLRSC can be improved:

- Running time and memory requirements. Currently cLRSC increases the running time and memory requirements for subspace clustering when compared to LRR.
- Strict enforcement of the rank constraint.

Recent advances in the field of Riemannian optimization [51] show promise in addressing both issues.

Acknowledgments

The research project is supported by the Australian Research Council (ARC) through grant DP140102270. It is also partly supported by the Natural Science Foundation of China (NSFC) through project 41371362.

References

[1] R. Bellman, R. Corporation, Dynamic Programming, Rand Corporation research study, Princeton University Press, 1957.
URL https://books.google.com.au/books?id=wdtoPwAACAAJ

[2] J. Friedman, T. Hastie, R. Tibshirani, The elements of statistical learning, Vol. 1, Springer series in statistics Springer, Berlin, 2001.

[3] I. Jolliffe, Principal Component Analysis, Wiley Online Library, 2002.

[4] J. Ramsay, B. W. Silverman, Functional Data Analysis, Springer Series in Statistics, Springer, 2005.

[5] J. Jacques, C. Preda, Functional data clustering: a survey, Advances in Data Analysis and Classification 8 (3) (2014) 231–255.

[6] C. Tomasi, T. Kanade, Shape and motion from image streams under orthography: a factorization method, International Journal of Computer Vision 9 (2) (1992) 137–154.

[7] J. P. Costeira, T. Kanade, A multibody factorization method for independently moving objects, International Journal of Computer Vision 29 (3) (1998) 159–179.

[8] K. Kanatant, Motion segmentation by subspace separation: Model selection and reliability evaluation, International Journal of Image and Graphics 2 (2) (2002) 179–197.

[9] B. Jacquet, R. Angst, M. Pollefeys, Articulated and restricted motion subspaces and their signatures, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013, pp. 1506–1513.

[10] R. Basri, D. W. Jacobs, Lambertian reflectance and linear subspaces, IEEE Transactions on Pattern Analysis and Machine Intelligence 25 (2) (2003) 218–233.

[11] A. S. Georghiades, P. N. Belhumeur, D. J. Kriegman, From few to many: Illumination cone models for face recognition under variable lighting and pose, IEEE Transactions on Pattern Analysis and Machine Intelligence 23 (6) (2001) 643–660.

[12] E. Elhamifar, R. Vidal, Sparse subspace clustering: Algorithm, theory, and applications, IEEE Transactions on Pattern Analysis and Machine Intelligence.

[13] W. Hong, J. Wright, K. Huang, Y. Ma, Multiscale hybrid linear models for lossy image representation, IEEE Transactions on Image Processing 15 (12) (2006) 3655–3671.

[14] Y. Zhang, Z. Jiang, L. S. Davis, Learning structured low-rank representations for image classification, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013, pp. 676–683.

[15] G. Bull, J. Gao, Transposed low rank representation for image classification, in: International Conference on Digital Image Computing Techniques and Applications, DICTA, Fremantle, Australia, 2012, pp. 1–7. doi:10.1109/DICTA.2012.6411718.

[16] R. Liu, Z. Lin, F. De la Torre, Z. Su, Fixed-rank representation for unsupervised visual learning, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 598–605.

[17] G. Liu, S. Yan, Latent low-rank representation for subspace segmentation and feature extraction, in: IEEE International Conference on Computer Vision (ICCV), 2011, pp. 1615–1622.

[18] A. Y. Yang, J. Wright, Y. Ma, S. S. Sastry, Unsupervised segmentation of natural images via lossy data compression, Computer Vision and Image Understanding 110 (2) (2008) 212–225.

[19] B. Cheng, G. Liu, J. Wang, Z. Huang, S. Yan, Multi-task low-rank affinity pursuit for image segmentation, in: IEEE International Conference on Computer Vision (ICCV), 2011, pp. 2439–2446.

[20] T. Hastie, P. Y. Simard, Metrics and models for handwritten character recognition, Statistical Science (1998) 54–65.

[21] L. Parsons, E. Haque, H. Liu, Subspace clustering for high dimensional data: a review, ACM SIGKDD Explorations Newsletter 6 (1) (2004) 90–105.

[22] T. Wallace, A. Sekmen, X. Wang, Application of subspace clustering in dna sequence analysis, Journal of Computational Biology 22 (10) (2015) 940–952.

[23] H. Wang, F. Chu, W. Fan, P. S. Yu, J. Pei, A fast algorithm for subspace clustering by pattern similarity, in: Scientific and Statistical Database Management, 2004. Proceedings. 16th International Conference on, IEEE, 2004, pp. 51–60.

[24] A. B. Tchagang, F. Famili, Y. Pan, Subspace clustering of dna microarray data: Theory, evaluation, and applications, International Journal of Computational Models and Algorithms in Medicine (IJCMAM) 4 (2) (2014) 1–52.

[25] F. Parvaresh, H. Vikalo, S. Misra, B. Hassibi, Recovering sparse signals using sparse measurement matrices in compressed dna microarrays, IEEE Journal of Selected Topics in Signal Processing 2 (3) (2008) 275–285.

[26] Y. Zhu, D. Huang, F. D. L. Torre, S. Lucey, Complex non-rigid motion 3d reconstruction by union of subspaces, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 1542–1549.

[27] A. Y. Yang, S. Iyengar, P. Kuryloski, R. Jafari, Distributed segmentation and classification of human actions using a wearable motion sensor

network, in: Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on, IEEE, 2008, pp. 1–8.

[28] S. Li, K. Li, Y. Fu, Temporal subspace clustering for human motion segmentation, in: Proceedings of the IEEE International Conference on Computer Vision, 2015, pp. 4453–4461.

[29] E. J. Candès, X. Li, Y. Ma, J. Wright, Robust principal component analysis?, Journal of the ACM 58 (3) (2011) 11.

[30] U. Von Luxburg, A tutorial on spectral clustering, Statistics and computing 17 (4) (2007) 395–416.

[31] J. Shi, J. Malik, Normalized cuts and image segmentation, IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (8) (2000) 888–905.

[32] C. Fowlkes, S. Belongie, F. Chung, J. Malik, Spectral grouping using the Nystrom method, IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (2) (2004) 214–225.

[33] E. Elhamifar, R. Vidal, Sparse subspace clustering, in: IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 2790–2797.

[34] M. Soltanolkotabi, E. Elhamifar, E. J. Candes, et al., Robust subspace clustering, The Annals of Statistics 42 (2) (2014) 669–699.

[35] G. Liu, Z. Lin, Y. Yu, Robust subspace segmentation by low-rank representation, in: International Conference on Machine Learning, 2010, pp. 663–670.

[36] W. Siming, L. Zhouchen, Analysis and improvement of low rank representation for subspace segmentation, arXiv preprint arXiv:1107.1561.

[37] R. Vidal, P. Favaro, Low rank subspace clustering (LRSC), Pattern Recognition Letters 43 (0) (2014) 47 – 61, {ICPR2012} Awarded Papers. doi:http://dx.doi.org/10.1016/j.patrec.2013.08.006.

[38] E. Liberty, F. Woolfe, P.-G. Martinsson, V. Rokhlin, M. Tygert, Randomized algorithms for the low-rank approximation of matrices, Proceedings of the National Academy of Sciences 104 (51) (2007) 20167–20172.

[39] F. Woolfe, E. Liberty, V. Rokhlin, M. Tygert, A fast randomized algorithm for the approximation of matrices, Applied and Computational Harmonic Analysis 25 (3) (2008) 335–366.

[40] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, Y. Ma, Robust recovery of subspace structures by low-rank representation, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (1) (2013) 171–184. doi:10.1109/TPAMI.2012.88.

[41] Q. Chen, Z. Shi, Z. An, Hyperspectral image fusion based on sparse constraint {NMF}, Optik - International Journal for Light and Electron Optics 125 (2) (2014) 832 – 838. doi:http://dx.doi.org/10.1016/j.ijleo.2013.07.061.
URL http://www.sciencedirect.com/science/article/pii/S0030402613010425

[42] S. Tierney, J. Gao, Y. Guo, Affinity pansharpening and image fusion, in: 2014 International Conference on Digital Image Computing: Techniques and Applications, DICTA 2014, Wollongong, New South Wales, Australia, November 25-27, 2014, 2014, pp. 1–8. doi:10.1109/DICTA.2014.7008094.
URL http://dx.doi.org/10.1109/DICTA.2014.7008094

[43] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, A. Blake, Real-time human pose recognition in parts from single depth images, in: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, 2011, pp. 1297–1304. doi:10.1109/CVPR.2011.5995316.

[44] J. Taylor, J. Shotton, T. Sharp, A. Fitzgibbon, The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation, in:

Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 103–110.

[45] L.-J. Li, H. Su, L. Fei-Fei, E. P. Xing, Object bank: A high-level image representation for scene classification & semantic feature sparsification, in: Advances in neural information processing systems, 2010, pp. 1378–1386.

[46] Q. V. Le, Building high-level features using large scale unsupervised learning, in: Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on, IEEE, 2013, pp. 8595–8598.

[47] R. Cabral, F. D. L. Torre, J. P. Costeira, A. Bernardino, Unifying nuclear norm and bilinear factorization approaches for low-rank matrix decomposition, in: Computer Vision (ICCV), 2013 IEEE International Conference on, IEEE, 2013, pp. 2488–2495.

[48] R. Liu, Z. Lin, Z. Su, Linearized alternating direction method with parallel splitting and adaptive penalty for separable convex programs in machine learning, in: ACML, 2013, pp. 116–132.

[49] N. Parikh, S. Boyd, Proximal algorithms, Foundations and Trends in Optimization 1 (3) (2013) 123–231.

[50] J.-F. Cai, E. J. Candès, Z. Shen, A singular value thresholding algorithm for matrix completion, SIAM Journal on Optimization 20 (4) (2010) 1956–1982.

[51] B. Vandereycken, Low-rank matrix completion by riemannian optimization, SIAM Journal on Optimization 23 (2) (2013) 1214—1236.